HashiCorp Consul

Unlocking the Cloud Operating Model: Networking

Contents

Overview

For most enterprises, digital transformation efforts mean delivering new business and customer value more quickly, and at a very large scale. The implication for Enterprise IT then is a shift from cost optimization to speed optimization. The cloud is an inevitable part of this shift as it presents the opportunity to rapidly deploy on-demand services with limitless scale.

To unlock the fastest path to value in the cloud, enterprises must consider how to industrialize the application delivery process across each layer of the cloud: embracing the cloud operating model, and tuning people, process, and tools to it.

In this white paper, we look at the implications of the cloud operating model, and present solutions for IT teams to adopt this model at the networking layer.

Challenges with multi-cloud service networking

In a traditional static network, monolithic applications are connected and secured by a fleet of network middleware, such as load balancers and firewalls. Companies focus on managing the traffic flowing in and out of the private datacenter (also known as north-south traffic). The network is segmented by coarse-grained network segments, such as staging, production, and PCI zones. This means it could take days or weeks to update a chain of network middleware in order to launch a new application.

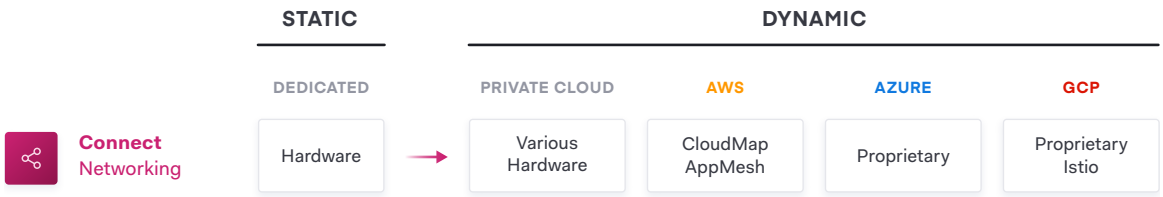

Figure 1: Static to Dynamic infrastructure in the network layer

The trend of application modernization, compounded by cloud adoption has shifted the underlying networking layer from host-based static networking to service-based dynamic networking.

In today's multi-cloud network:

- Micro services are running on ephemeral containers with dynamic IPs.
- The explosion of service-to-service communication (also known as east-west traffic) result in complex traffic patterns.
- The public clouds blur the boundary of network perimeters, presenting a new set of security challenges.

As the network becomes larger and more dynamic, it puts pressure on networking and security teams to keep pace with more changes to existing network middleware. These teams often cannot keep up with the increasing demand, and their manual approach to network operation undermines the agility of application teams.

HashiCorp Consul: Multi-cloud service networking simplified

Designed to solve the rising networking challenges for dynamic infrastructure, Consul provides a multi-cloud networking platform to connect and secure services.

Consul takes a different approach to networking by:

- Elevating the operation from IP address to service level.
- Providing a central service registry to locate the dynamic location of services.
- Enabling service discovery to allow services to register, discover, and connect with each other directly.
- Driving automation to eliminate the operational burden of updating network middleware.
- Enacting a service mesh solution to simplify networking by shifting naming abstraction, routing, authorization, and other networking functionalities from centralized middleware to the endpoints.
- Being platform-agnostic, which allows IT teams to maintain a single networking workflow across different runtime platforms, public clouds, and private clouds.

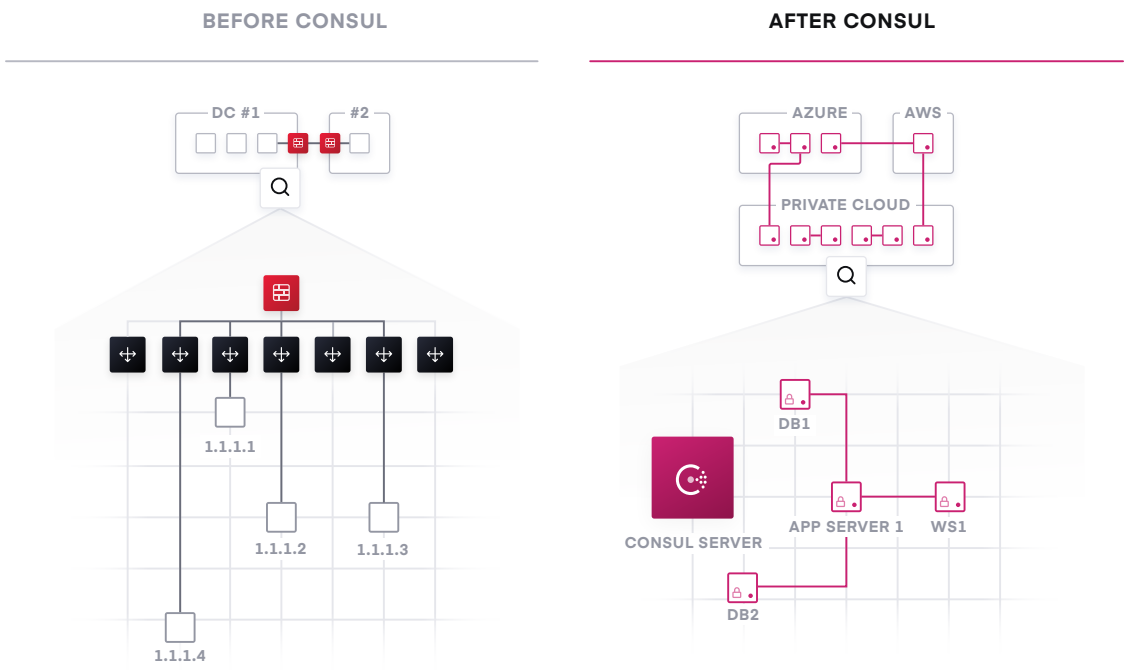

Figure 2: From traditional networking to modern service networking with Consul

Consul's crawl, walk, run journey

Adoption of HashiCorp Consul typically follows a three-stage pattern for most organizations. Each stage gradually adds more automation to the process and more software-driven technologies to improve agility, performance, and security and reduce cost.

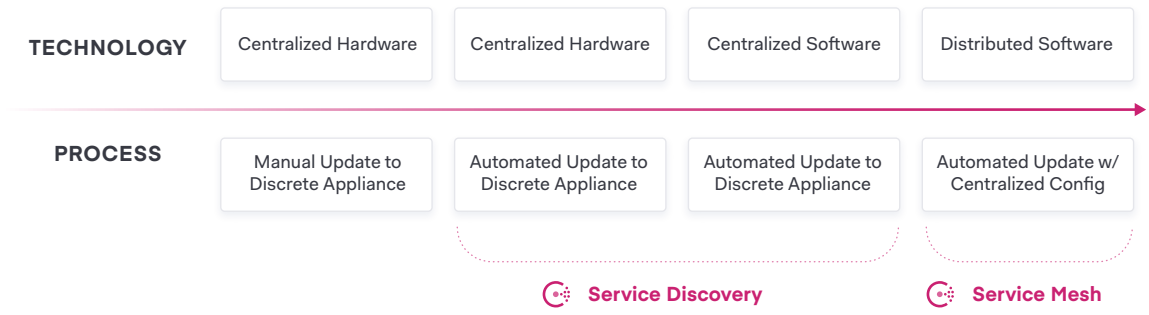

Figure 3: Evolution of technology and process in the network layer

Crawl: Central service registry

For organizations at the beginning of their cloud journey, finding a way to track and manage the explosion of services across multiple subnets, data centers, and cloud regions is their first hurdle.

The traditional approach relies on load balancers and virtual IPs to provide a naming abstraction to represent a service with a static IP. The process to track the network location of services often takes the form of spreadsheets, load balancer dashboards, or configuration files, all of which are disjointed, manual processes that are not ideal.

The starting point for networking in the new cloud operating model is the deployment of a common service registry, which provides a real-time "directory" of what services are running, where they are, and their current health status.

A lightweight Consul agent runs on every service, allowing each service instance to programmatically register with the registry as well as discover and connect with other services automatically through DNS or API interfaces. Additional tooling allows Consul to register and monitor appliances and devices that cannot run an agent.

The integrated health check will monitor each service instance's health status so the IT team can triage the availability of each instance and Consul can help prevent routing traffic to unhealthy service instances.

Walk: Network automation with service discovery

The next step is to unlock agility with existing networking architecture through network automation. Instead of a manual, ticket-based process to reconfigure the traditional network middleware every time there is a change in service network locations or configurations, Consul provides a publisher/subscriber (PubSub) service to automate these network operations.

The newly added or removed service instances will automatically "publish" their location information with the service registry. The network middleware can subscribe to service changes from the service registry (by using tools like Consul Template or native integration). This enables a publish/subscribe style of automation that can handle highly dynamic infrastructure and scale infinitely.

For example, F5 BIG-IP devices support native Consul integration. Instead of being configured to balance traffic to a static pool of backend services, the BIG-IP can be configured to route to a dynamic set of backends that are provided by Consul. The moment an additional backend service is provisioned and registered, the BIG-IP will automatically update the set of backends and begin routing traffic to it.

Traffic can be routed to new service instances instantly without manual intervention. The same approach applies to firewalls and other critical network middleware as well. This publish/subscribe model significantly accelerates the productivity of the network and IT operation teams by decoupling their workflow from IP addresses. It also decouples the workflow between teams, as operators can independently deploy applications and publish to Consul, while operations teams can subscribe to handle downstream automation.

Run: Service mesh

As organizations continue to scale with microservices-based or cloud-native applications, the underlying infrastructure becomes larger and more dynamic. Modular services need to communicate with each other to compose business functionality, leading to an explosion of east-west traffic.

Existing networking approaches with network appliances cannot effectively handle east-west traffic in dynamic settings. They cause a proliferation of expensive network middleware, introduce single points of failure all over the system, and add significant operational overhead to IT teams.

A distributed service mesh pushes routing, authorization and other networking functionalities to the endpoints in the network, rather than imposing them through middleware. This makes the network topology simpler and easier to manage, it removes the need for expensive middleware within east-west traffic paths, and it makes service-to-service communication much more reliable and scalable because of the network's decentralized nature.

Consul provides an API-driven control plane, which integrates with sidecar proxies alongside each service instance (proxies such as Envoy, HAProxy, and NGINX) that provide the distributed data plane.

The service mesh approach allows critical functionality like naming, segmentation, authorization, traffic management, and observability to be configured through policies in the central registry and enforced by proxies at the endpoint where an individual service is running.

Consul enables a zero-trust network model by securing service-to-service communication with automatic TLS encryption and identity-based authorization. Network operation and security teams can define the security policies through *intentions* with logical services rather than IP addresses. For example, allowing web servers to communicate with databases, instead of IP1 to IP2. Proxies will enforce security consistently regardless of how services scale up and down or migrate to other platforms.

Consul can also be integrated with HashiCorp Vault for centralized PKI and certificate management.

The security benefits of a mesh based approach are several fold. For most organizations, traffic within a network zone (such as production or PCI) is relatively flat. This means a compromise of a single service would allow an attacker to move laterally to other systems in the same zone. Consul enables a much more fine grained authorization of access to avoid this.

By handling segmentation above the network layer, Consul avoids the need to have complex cloud topologies with hundreds or thousands of accounts. Instead a few accounts can be used, and the logical segmentation is done through the use of proxies and centrally defined rules.

Summary

In a multi-cloud data center, networking services should be provided centrally, whereby IT teams starts with providing service registry and service discovery capabilities. Having a common registry provides a "map" of what services are running, where they are, and their current health status.

The registry can be queried programmatically to drive network automation of API gateways, load balancers, firewalls, and other critical middleware components used for north-south traffic. And network middleware components used for east-west traffic can be replaced by using Consul service mesh, where proxies run on the edge to provide the same functionality with greater scalability and security.

Consul's crawl, walk, run journey is a simple path that organizations can follow towards a more simplified and scalable networking approach to support distributed applications with multi-cloud and multi-data center topologies.


HashiCorp